Inapproximability after Uniqueness Phase Transition in Two-Spin Systems

Jin-Yi Cai[1,⋆], Xi Chen[2,⋆⋆], Heng Guo[1,⋆⋆⋆], and Pinyan Lu[3]

[1] University of Wisconsin, Madison
[2] Columbia University
[3] Microsoft Research Asia

Abstract. A two-state spin system is specified by a matrix

$$\mathbf{A} = \begin{bmatrix} A_{0,0} & A_{0,1} \\ A_{1,0} & A_{1,1} \end{bmatrix} = \begin{bmatrix} \beta & 1 \\ 1 & \gamma \end{bmatrix} \tag{1}$$

where $\beta, \gamma \geq 0$. Given an input graph $G = (V, E)$, the partition function $Z_{\mathbf{A}}(G)$ of a system is defined as

$$Z_{\mathbf{A}}(G) = \sum_{\sigma: V \to \{0,1\}} \prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}. \tag{2}$$

We prove inapproximability results for the partition function $Z_{\mathbf{A}}(G)$ in the region specified by the non-uniqueness condition from phase transition for the Gibbs measure. More specifically, assuming NP $\neq$ RP, for any fixed β, γ in the unit square, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A}}(G)$ for d-regular graphs G with relative error $\epsilon = 10^{-4}$, if $d = \Omega(\Delta(\beta, \gamma))$, where $\Delta(\beta, \gamma) > 1/(1 - \beta\gamma)$ is the uniqueness threshold. Up to a constant factor, this hardness result confirms the conjecture that the uniqueness phase transition coincides with the transition from computational tractability to intractability for $Z_{\mathbf{A}}(G)$. We also show a matching inapproximability result for a region of parameters β, γ outside the unit square, and all our results generalize to partition functions with an external field.

1 Introduction

Spin systems are well studied in statistical physics and applied probability. We focus on two-state spin systems. An instance of a spin system is a graph $G = (V, E)$. A configuration $\sigma : V \to \{0, 1\}$ assigns to each vertex one of two states. The contributions of local interactions between adjacent vertices are quantified by (1), a 2×2 matrix with $\beta, \gamma \geq 0$. The partition function $Z_{\mathbf{A}}(G)$ of a system is defined by (2), and we use $\omega(G, \sigma)$ to denote the weight of σ:

$$\omega(G, \sigma) = \textstyle\prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}$$

⋆ Supported by NSF CCF-0914969.
⋆⋆ Supported by NSF CCF-1139915 and start-up funds from Columbia University.
⋆⋆⋆ Supported by NSF CCF-0914969.

G. Lin (Ed.): COCOA 2012, LNCS 7402, pp. 336–347, 2012.

Given a fixed $\mathbf{A}$, we are interested in the complexity of computing $Z_{\mathbf{A}}(G)$, where G is given as an input. Many natural combinatorial counting problems can be formulated as two-state spin systems. For example, with $(\beta, \gamma) = (0, 1)$, $Z_{\mathbf{A}}(G)$ is exactly the number of independent sets (or vertex covers) of G. The definition of $Z_{\mathbf{A}}(G)$ in (2) can also be generalized to larger matrices $\mathbf{A}$, and the problem is known as counting (weighted) graph homomorphisms [1, 2]. On the other hand, the so-called Ising model is the special case where $\beta = \gamma$.

The *exact* complexity of computing $Z_{\mathbf{A}}(G)$ has been completely solved for any fixed symmetric $\mathbf{A}$ [3–6] and even for not necessarily symmetric $\mathbf{A}$ [7–12] as part of the dichotomy theorems for general counting constraint satisfaction problems (#CSP). When specialized to two-state spin systems, $Z_{\mathbf{A}}(G)$ is #P-hard to compute exactly, except for the two restricted settings of $\beta\gamma = 1$ or $\beta = \gamma = 0$ for which cases $Z_{\mathbf{A}}(G)$ is polynomial-time computable. Consequently, the study on two-state spin systems has focused on the approximation of $Z_{\mathbf{A}}(G)$, and this is the subject of the present paper.

Following standard definitions, a fully polynomial-time approximation scheme (FPTAS) for $Z_{\mathbf{A}}(G)$ is an algorithm that, given as input a graph G as well as a parameter $\epsilon > 0$, outputs a number Z that satisfies

$$(1 - \epsilon) \cdot Z_{\mathbf{A}}(G) \leq Z \leq (1 + \epsilon) \cdot Z_{\mathbf{A}}(G) \tag{3}$$

in time poly$(|G|, 1/\epsilon)$. A fully polynomial-time randomized approximation scheme (FPRAS) is a randomized algorithm that, with probability $1 - \delta$, outputs a number Z satisfying (3) in time poly$(|G|, 1/\epsilon, \log(1/\delta))$.

In a seminal paper [13] Jerrum and Sinclair gave an FPRAS for $Z_{\mathbf{A}}(\cdot)$ with $\beta = \gamma > 1$. It was then further extended to the entire region of $\beta\gamma > 1$ by Goldberg, Jerrum and Paterson [14]. We call a two-state spin system *ferromagnetic* if $\beta\gamma > 1$ and *anti-ferromagnetic* if $\beta\gamma < 1$. The approximability of $Z_{\mathbf{A}}(\cdot)$ for anti-ferromagnetic systems is less well understood. Starting with counting independent sets in sparse graphs [15], the approximability of $Z_{\mathbf{A}}(\cdot)$ in bounded degree graphs is also widely studied. Significant progress has been made recently on the algorithmic side, and approximation algorithms for anti-ferromagnetic two-state spin systems have been developed [16–19], based on the technique of correlation decay introduced by Bandyopadhyay and Gamarnik [20] and Weitz [16]. Finally a unified FPTAS was found [19] to approximate $Z_{\mathbf{A}}(\cdot)$ for all anti-ferromagnetic two-state spin systems of either bounded degree graphs or general graphs, when the system satisfies a *uniqueness condition*.

The uniqueness condition is named for, and closely related to, phase transitions that occur for the Gibbs measure. It depends on not only β, γ but also the degree of the underlying graph. Such phase transitions from statistical physics are believed to frequently coincide with the transitions of computational complexity from tractability to intractability. However, there are only very few examples where this conjectured link is rigorously proved. One notable example is for the hardcore gas model (or independent set, with $\beta = 0$ and $\gamma = 1$), for which such a conjecture was rigorously proved (for almost all degree bounds) both for the algorithmic side [16] and for the hardness side [21, 22]. As discussed

above [16–19], for general anti-ferromagnetic two-state spin systems, the algorithmic part of the conjecture has recently been established. In this paper, we make substantial progress on the hardness part of the conjecture.

Our Results. For $0 \leq \beta, \gamma \leq 1$ except at $(\beta, \gamma) = (0, 0)$ and $(1, 1)$, Goldberg, Jerrum and Paterson proved that the problem does not admit an FPRAS for general graphs (when there is no degree bound), unless NP = RP [14]. In their reduction, the degrees of the hard instances are unbounded. This is consistent with the uniqueness threshold conjecture. However, for any fixed β and γ in the unit square, the uniqueness condition states that there exists a finite threshold degree $\Delta(\beta, \gamma)$ [17–19], which satisfies

$$\Delta(\beta, \gamma) > \frac{1 + \sqrt{\beta\gamma}}{1 - \sqrt{\beta\gamma}} = \frac{(1 + \sqrt{\beta\gamma})^2}{1 - \beta\gamma} \geq \frac{1}{1 - \beta\gamma} \tag{4}$$

such that the system satisfies the uniqueness condition if the degree $d < \Delta(\beta, \gamma)$ and the non-uniqueness condition if $d \geq \Delta(\beta, \gamma)$. The paper [19] gives an FPTAS for graphs with degree bounded by $\Delta(\beta, \gamma)$. The conjectured coincidence of phase transition with hardness in complexity suggests that as soon as the degree of the input graph goes beyond $\Delta(\beta, \gamma)$, $Z_{\mathbf{A}}(\cdot)$ becomes hard to approximate. Towards this direction, we show that for any fixed β, γ in the unit square, the problem does not have an FPRAS if the degree of the input graph is $\Omega(\Delta(\beta, \gamma))$, unless NP = RP. Our hardness result also holds when restricted to input graphs that are regular. Formally, we prove the following theorem:

Theorem 1. *There exists a positive constant h with the following property. For any $\beta, \gamma : 0 \leq \beta, \gamma \leq 1$ such that $(\beta, \gamma) \neq (0, 0)$, $(1, 1)$ and for any integer $d \geq h/(1 - \beta\gamma)$, there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A}}(G)$ in d-regular graphs G with relative error $\epsilon = 10^{-4}$, unless* NP = RP.

Note the relation between our degree bound $h/(1 - \beta\gamma)$ and $\Delta(\beta, \gamma)$ from (4).

We also make progress on β, γ outside the unit square. While the uniqueness condition is *monotone* inside the unit square, its behavior outside is significantly different. (See more discussions on this difference in the appendix of the full version [23].) Without loss of generality, we consider the region defined by $\beta\gamma < 1$ with $0 < \beta < 1 < \gamma$. There is a uniqueness curve (see Figure 1), connecting the point $(1, 1)$ and the γ-axis. Above the curve, the system satisfies the uniqueness condition for any graph [18, 19]. Hence, hardness is only possible below the uniqueness curve. Furthermore, when (β, γ) is outside the unit square but below this uniqueness curve, there is only a finite range of degrees d for which the system does not satisfy the uniqueness condition. This makes it very challenging to prove a hardness result for them. Previously, the hardness was only obtained in [14] for a very tiny square $0 \leq \beta \leq \eta$ and $1 \leq \gamma \leq 1 + \eta$ where η is roughly 10^{-7}, near $(0, 1)$ corresponding to the hardcore gas model (independent set).

We prove the following hardness result for (β, γ) outside the unit square:

Theorem 2. *Given β and γ such that $0 < \beta < 1$, $\gamma > 1$ and $\beta\gamma < 1$, let*

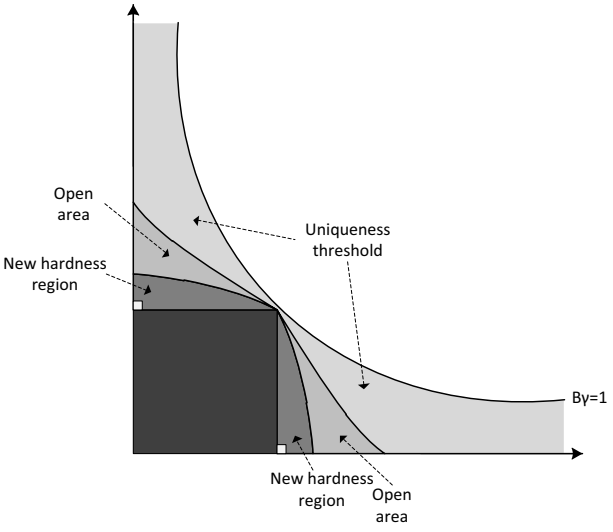

Fig. 1. The new hardness region of Theorem 2

$$\Delta' = \left\lceil -1/(\ln\beta + \ln\gamma) \right\rceil \quad \textit{and} \quad \Delta^* = \left\lceil 1/\ln\gamma \right\rceil . \tag{5}$$

When $\Delta^* \geq 8000\Delta'$*, there exists no randomized polynomial-time algorithm that approximates* $Z_{\mathbf{A}}(G)$ *in regular graphs of degree* Δ^* *with relative error* $\epsilon = 10^{-4}$*, unless* NP = RP.

The new hardness region is pictured in Figure 1 [1] above. Here the two white squares are the hardness regions acquired by Goldberg, Jerrum, and Paterson [14]. Beyond the uniqueness threshold, we know that FPTAS exists. Our hardness result, Theorem 2, applies to the region between the vertical line with $\gamma = 1$ and the curve to the left of the uniqueness threshold. Let us describe this new curve in more details. We focus on the region with $0 < \beta < 1 < \gamma$ and $\beta\gamma < 1$. There is a symmetric curve for $0 < \gamma < 1 < \beta$. Near the point $(1, 1)$, the condition imposed by Theorem 2 is almost linear. So the new curve is roughly a line with slope -8000 around $(1, 1)$. When it approaches the line of $\beta = 0$, Δ' becomes 1 and the condition requires γ to be between 1 and roughly $1 + 1/8000$.

Moreover, using a standard translation (see the appendix of the full version [23]), we can generalize both Theorem 1 and 2 to two-state spin systems with an external field. Formally, let $\mu \geq 0$, we have the following two corollaries for

$$Z_{\mathbf{A},\mu}(G) = \sum_{\sigma: V \to \{0,1\}} \mu^{\left|\{v \in V : \sigma(v) = 0\}\right|} \prod_{(u,v) \in E} A_{\sigma(u),\sigma(v)}.$$

[1] The reader should be aware that, for illustration purposes, the picture is not drawn to actual scale.

Corollary 1. *There is a positive constant h with the following property. Given any non-negative β, γ and μ with $\beta\gamma < 1$, if d is a positive integer that satisfies $\gamma \le \mu^{\frac{1}{d}} \le 1/\beta$ and $d \ge h/(1-\beta\gamma)$, then there exists no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A},\mu}(G)$ in d-regular graphs with relative error $\epsilon = 10^{-4}$, unless* NP = RP.

Corollary 2. *Given any non-negative β, γ, μ and a positive integer d such that*

$$e^{\frac{1}{d}} \le \gamma \cdot \mu^{-\frac{1}{d}} < e^{\frac{1}{d-1}} \quad \text{or} \quad e^{\frac{1}{d}} \le \beta \cdot \mu^{\frac{1}{d}} < e^{\frac{1}{d-1}},$$

if d also satisfies $d \ge 8000\,\lceil -1/(\ln\beta + \ln\gamma)\rceil$, then there is no randomized polynomial-time algorithm that approximates $Z_{\mathbf{A},\mu}(G)$ in d-regular graphs with relative error $\epsilon = 10^{-4}$, unless NP = RP.

Proof Outline. In the proofs of both Theorem 1 and Theorem 2, we use the phase transition that occurs in the non-uniqueness region to encode a hard-to-approximate problem. This approach has been used in previous hardness proofs for the hardcore gas model [15, 21, 24]. To this end, we reduce the approximation of E2LIN2 to the approximation of partition function in a two-state spin system. Here an instance of E2LIN2 consists of a set of variables $x_1, \ldots, x_n$ and a set of equations of the form $x_i + x_j = 0$ or 1 over $\mathbb{Z}_2$. From [25], it is NP-hard to approximate the number of satisfiable equations for E2LIN2 within any constant factor better than 11/12.

Given an E2LIN2 instance with variables $x_1, \ldots, x_n$, we use a random bipartite regular graph to encode each variable x_i. Due to the phase transition and the fact that we are in a non-uniqueness region, each of these bipartite regular graphs would be in one of two types of configurations with high probability, if sampled proportional to its weight in the partition function. This can be used to establish a correspondence between the configurations of these bipartite graphs and the assignment of the n boolean variables $x_1, \ldots, x_n$. Furthermore, we also add external connections between the random bipartite graphs according to the set of equations in the E2LIN2 instance. They contribute exponentially to the total weight in the partition function, according to the total number of equations that an assignment satisfies. Thus, a sufficiently good approximation to the partition function can be used to decode approximately the maximum number of equations that an assignment can satisfy.

Our gadget is also randomly constructed, so the probability should also be over the distribution of the gadgets. It is not hard to show that things work out beautifully if we simply substitute the expectation for the actual weight. But to make the proof rigorous, one must first obtain a sufficiently good concentration result. Such a result is unknown and could be very difficult to prove (assuming it is true), as it is already a tour-de-force in the special case for the hardcore gas model [21, 22, 24].

Instead, we use a detour: (1) We prove a lower bound for the weights of the two types of configurations we expect, guided by the phase transition; and (2) We prove that the total weight of other configurations is exponentially smaller compared to the lower bound proved in (1), with probability exponentially close

to 1. The way we establish the lower bound in (1) is similar to the approach of Dyer, Frieze, and Jerrum [15]. To prove (2), they [15] used the expectation and Markov's inequality. If we use the same approach, we could not get the hardness result for bounded degree graphs in the same order of the uniqueness bound. Instead, we use a new approach for (2).

Indeed we first show a high concentration result for an expander property of the gadgets we use. Then we show that the total weight of other configurations must be exponentially small, given that the gadgets satisfy that property. This circumvented our inability to prove a complete concentration result. But we do need to prove some limited concentration results regarding the random gadget. This then led us to the hardness results for degrees in the right order conjectured according to the uniqueness threshold. It remains open whether one can use a refined version of this reduction along with the proof by Sly [21] to get the exact right bound. As discussed in the appendix of the full version [23], this random regular graph follows quite closely the property of phase transition in infinite d-ary trees, when the parameter is below or beyond the uniqueness condition.

While the high-level idea of our proofs for both Theorem 1 and Theorem 2 are quite clear and similar, it remains a challenge to work out the estimation for all ranges of parameters and at the same time, make sure that the degree is in the same order of the uniqueness bound. To this end, technically we need to use very different approaches for Theorem 1 and Theorem 2. Even within Theorem 1 itself, we need to do the estimation differently for three different subcases.

2 Proof of the Main Theorems

From now on, we will use $Z(G)$ to denote $Z_{\mathbf{A}}(G)$ whenever it is clear from the context. Given positive integers N and Δ, we use $\mathcal{H}(N, \Delta)$ to denote the following probability distribution of Δ-regular bipartite graphs $H = (U \cup V, E)$ with bipartition U, V and $|U| = |V| = N$: H is the union of Δ perfect matchings between U and V each selected independently and uniformly at random. (Because these perfect matchings are drawn independently, H may have parallel edges.)

In the proofs of both Theorem 1 and 2, we give a polynomial-time reduction from E2LIN2 to the approximation of $Z(G)$. An instance of E2LIN2 consists of m equations over $\mathbb{Z}_2$ in n variables $x_1, \ldots, x_n$. Each equation has exactly two variables and is of the form $x_i + x_j = b \in \{0, 1\}$. Without loss of generality we may always assume $m \geq n/2$; otherwise one of the variables does not appear in any equation. Given an assignment S of the n variables $x_1, \ldots, x_n$, we use $\theta(S)$ to denote the number of equations that S satisfies and let $\theta^* = \max_S \theta(S)$. In [25] Håstad showed that it is NP-hard to estimate θ^* within any constant factor better than 11/12.

Given an E2LIN2 instance, we construct a random $(\Delta + \Delta')$-regular graph G as follows, with the two parameters Δ, Δ' to be specified later. This construction is used in the proofs of both Theorem 1 and 2:

Construction of G from an instance of E2LIN2. For each variable x_i, $i \in [n]$, we let U_i and V_i denote two sets of $d_i m$ vertices each, where $d_i \geq 1$

denotes the number of equations in which x_i appears (thus, $\sum_i d_i = 2m$). Moreover, U_i and V_i can be decomposed into

$$U_i = U_{i,1} \cup \cdots \cup U_{i,d_i} \quad \text{and} \quad V_i = V_{i,1} \cup \cdots \cup V_{i,d_i}$$

where each $U_{i,k}$ and $V_{i,k}$ contains exactly m vertices. Now enumerate all the m equations in the E2LIN2 instance one by one. For each of the m equations do the following:

(1) Let $x_i + x_j = b \in \{0,1\}$ denote the current equation. Assume this is the kth time that x_i appears in an equation, and the ℓth time that x_j appears in an equation so far, where $k \in [d_i]$ and $\ell \in [d_j]$. Denote the m vertices in $U_{i,k}$ by $\{u_1, \ldots, u_m\}$, vertices in $V_{i,k}$ by $\{v_1, \ldots, v_m\}$, vertices in $U_{j,\ell}$ by $\{u'_1, \ldots, u'_m\}$ and vertices in $V_{j,\ell}$ by $\{v'_1, \ldots, v'_m\}$. All these vertices have degree 0 at this moment. If $b = 0$, we add Δ' parallel edges between (u_s, v'_s) and (v_s, u'_s), for each $s \in [m]$; or if $b = 1$, we add Δ' parallel edges between (u_s, u'_s) and (v_s, v'_s), for each $s \in [m]$.

By the end of this step, every vertex has degree Δ'. In the next step,

(2) For each $i \in [n]$, we add a bipartite graph $H_i = (U_i \cup V_i, E_i)$ drawn from $\mathcal{H}(d_i m, \Delta)$.

This finishes the construction and we get a $(\Delta + \Delta')$-regular graph G with $4m^2$ vertices.

We need the following notation. Given an assignment $\sigma : V(G) \to \{0,1\}$, we use $U_i(\sigma)$ to denote the number of vertices $u \in U_i$ with $\sigma(u) = 0$, and use $V_i(\sigma)$ to denote the number of $v \in V_i$ with $\sigma(v) = 0$.

Proof (of Theorem 1). Without loss of generality, assume $\beta, \gamma : 0 \le \beta \le \gamma \le 1$. We can also assume $\beta > 0$, as the tight hardness to the exact uniqueness bound for $\beta = 0$ has been shown in [18], by generalizing the tight hardness result for the hardcore model [21, 22].

Given an assignment S of the n variables, we let $Z(G, S)$ denote the sum of $\omega(G, \sigma)$ over assignments $\sigma : V(G) \to \{0,1\}$ that satisfy for each $i \in [n]$,

$$U_i(\sigma) \le V_i(\sigma) \quad \text{if } x_i = 0 \text{ in } S; \text{ or } \quad U_i(\sigma) \ge V_i(\sigma) \quad \text{if } x_i = 1 \text{ in } S. \tag{6}$$

From definition we have $Z(G, S) \le Z(G) \le \sum_S Z(G, S)$. We need the following key lemma. Its proof can be found in the full version [23]:

Lemma 1. *There exists a positive constant h with the following property: For any β and $\gamma : 0 < \beta \le \gamma \le 1$ with $(\beta, \gamma) \ne (1,1)$ and for any $\Delta^* \ge h/(1 - \beta\gamma)$, there are $D > 1$, $C > 0$ and positive integers Δ and Δ' with $\Delta + \Delta' = \Delta^*$, that satisfy the following property: given any input instance of* E2LIN2 *with n variables $x_1, \ldots, x_n$ and m equations, except for probability $\le \exp(-\Omega(m))$, the Δ^*-regular graph G constructed with parameters Δ and Δ' satisfies*

$$C^{m^2} \cdot D^{m\theta(S)} \leq Z(G,S) \leq C^{m^2} \cdot D^{m\left(\theta(S)+0.03\,m\right)}, \tag{7}$$

for any assignment S of the n variables.

Given β, γ and Δ^*, we let C, D, Δ and Δ' denote the constants that satisfy the condition in Lemma 1. Then given an input instance of E2LIN2, (7) holds with probability $1 - \exp(-\Omega(m))$.

Now assume (7) holds. We use θ^* to denote the maximum number of consistent equations and use S^* to denote an assignment that satisfies θ^* equations. We also use Y to denote an estimate of $Z = Z(G)$, where $|Y/Z - 1| \leq \epsilon = 10^{-4}$. From (7) and $Z(G,S^*) \leq Z(G) \leq \sum_S Z(G,S)$, we get

$$(1+\epsilon) \cdot 2^n \cdot C^{m^2} \cdot D^{m(\theta^*+0.03\,m)} \geq Y \geq (1-\epsilon) \cdot C^{m^2} \cdot D^{m\theta^*} \tag{8}$$

Using Y, we set

$$Y' = \frac{\ln Y - \ln(1+\epsilon) - n\ln 2 - m^2 \ln C - 0.03\,m^2 \ln D}{m \ln D}$$

and we get $Y' \leq \theta^*$ since $\ln D > 0$. We finish the proof by showing that $Y' > (11/12) \cdot \theta^*$. By (8) we get

$$Y' \geq \theta^* - \frac{\ln(1+\epsilon) - \ln(1-\epsilon) + n\ln 2 + 0.03\,m^2 \ln D}{m \ln D}$$

As $\theta^* \geq m/2$ and $m \geq n/2$, when m is large enough, $Y' > (11/12) \cdot \theta^*$ and the theorem is proven.

Next, we prove Theorem 2:

Proof (of Theorem 2). For β, γ with $0 < \beta < 1 < \gamma$ and $\beta\gamma < 1$, let Δ' and Δ^* be the two positive integers defined in (5) which satisfy $\Delta^* \geq 8000\,\Delta'$. We set $\Delta = \Delta^* - \Delta'$. Given any input instance of E2LIN2 with n variables and m equations, we let G denote the Δ^*-regular graph constructed using Δ and Δ'.

First we show that to get a good approximation of $Z(G)$, with high probability it suffices to sum $\omega(G,\sigma)$ only over assignments σ satisfying the following:

$$\min\Big(U_i(\sigma), V_i(\sigma)\Big) \leq \lambda d_i m, \quad \text{for all } i \in [n], \text{ where } \lambda = 9 \times 10^{-5}. \tag{9}$$

We let Σ denote the set of all such assignments. Formally we prove the following key lemma in Section 3:

Lemma 2. *Let G be the graph constructed from an* E2LIN2 *instance with n variables $x_1, \ldots, x_n$ and m equations, with parameters Δ, Δ'. Then with probability $1 - \exp(-\Omega(m^{1/3}))$, it satisfies*

$$\sum_{\sigma \in \Sigma} \omega(G,\sigma) \leq Z(G) \leq \big(1+o(1)\big) \cdot \sum_{\sigma \in \Sigma} \omega(G,\sigma). \tag{10}$$

Next, given an assignment S over the n variables we use $Z_\Sigma(G,S)$ to denote the sum of $\omega(G,\sigma)$ over all assignments $\sigma \in \Sigma$ that satisfy (6) for all $i \in [n]$. We prove the following lemma in the full version [23]:

Lemma 3. *There are $C > 0$ and $D > 1$ satisfying the following property: given any instance of* E2LIN2 *with n variables and m equations, the Δ^*-regular graph G constructed with parameters Δ and Δ' satisfies*

$$C^{m^2} \cdot D^{m\theta(S)} \le Z_\Sigma(G,S) \le C^{m^2} \cdot D^{m\left(\theta(S)+0.04\,m\right)}, \tag{11}$$

for any assignment S of the n variables.

Let $\theta^* \ge m/2$ denote the maximum number of consistent equations, and let S^* denote an assignment that satisfies θ^* equations. From these two lemmas we have with high probability that

$$\begin{aligned} C^{m^2} \cdot D^{m\theta^*} \le Z_\Sigma(G,S^*) \le Z(G) &\le \big(1+o(1)\big) \cdot \sum\nolimits_S Z_\Sigma(G,S) \\ &\le \big(1+o(1)\big) \cdot 2^n \cdot C^{m^2} \cdot D^{m(\theta^*+0.04m)} \end{aligned}$$

Theorem 2 then follows from the same argument used in Theorem 1.

3 Proof of Lemma 2

Recall that β and γ satisfy $\beta,\gamma : 0 < \beta < 1 < \gamma$ and $\beta\gamma < 1$. Let Δ' and Δ^* be the two positive integers defined in Theorem 2, with $\Delta^* \ge 8000\Delta'$. From their definitions, we have $(\beta\gamma)^{\Delta'} \le 1/e$ and $\gamma^{\Delta^*} \ge e$. Set $\Delta = \Delta^* - \Delta' \ge 7999\Delta' \ge 7999$. By the definition of Δ^*, we have $e > \gamma^{\Delta^*-1} \ge \gamma^{\Delta}$ and thus, $\gamma < 1.001$.

Given an E2LIN2 instance with n variables $x_1,\ldots,x_n$ and m equations, we use G to denote the Δ^*-regular graph constructed with parameters Δ and Δ', where $\Delta^* = \Delta + \Delta'$. We let H_i denote the bipartite graph in G that corresponds to x_i and use $U_i \cup V_i$ to denote its vertices, with $|U_i| = |V_i| = d_i m$.

Before working on G and H_i, we start by proving a property that a bipartite graph sampled from the distribution $\mathcal{H}(N,\Delta)$ satisfies with very high probability. Let H be a bipartite graph drawn from $\mathcal{H}(N,\Delta)$ for some $N \ge 1$ and Δ defined above, with $2N$ vertices $U \cup V$. We also use $\rho : U \cup V \to \{0,1\}$ to denote an assignment and call it an (a,b)-assignment for some $a,b \in T_N$, where

$$T_N = \big\{0, 1/N, 2/N, \ldots, (N-1)/N, 1\big\}$$

if $|\,u \in U : \rho(u) = 0\,| = aN$ and $|\,v \in V : \rho(v) = 0\,| = bN$. We also use $\mathcal{I}_N(a,b)$, where $a,b \in T_N$, to denote the set of all such (a,b)-assignments, and let

$$Z_{a,b}(H) = \sum_{\rho \in \mathcal{I}_N(a,b)} \omega(H,\rho) \cdot \gamma^{\Delta'(2-a-b)N} \tag{12}$$

with Δ' as defined above. We are interested in the expectation of $Z_{a,b}(H)$ when $\min(a,b) \ge \lambda = 9 \times 10^{-5}$:

Lemma 4. *For large enough N and $a, b \in T_N$ such that* $\min(a, b) \geq \lambda$*, we have*

$$\mathbf{E}_{H \leftarrow \mathcal{H}(N,\Delta)}\Big[Z_{a,b}(H)\Big] \leq \exp\big(1.21 \cdot N\big).$$

The proof of Lemma 4 can be found in the full version [23]. By Lemma 4 we can impose the following condition on the graph G constructed from the input instance of E2LIN2: For all $i \in [n]$ and all $a, b \in T_{d_i m}$ with $\min(a, b) \geq \lambda$,

$$Z_{a,b}(H_i) \leq \exp\big(1.22 \cdot d_i m\big). \tag{13}$$

Using Lemma 4, Markov's inequality and the union bound, it is easy to show that G satisfies this condition with probability $1 - \exp(-\Omega(m))$.

In the rest of the proof, we prove that G satisfies (10) whenever it satisfies (13). Lemma 2 then follows immediately.

We assume that G satisfies (13). Then to prove (10), we randomly sample an assignment σ with probability proportional to $\omega(G, \sigma)$. (10) follows if we can show that σ satisfies (9) with probability $1 - o(1)$. For this purpose we need the following lemmas that give us properties that σ satisfies with high probability. Given any set L of vertices in G, we let

$$N_\sigma(L) = \big\{v \in L : \sigma(v) = 0\big\}.$$

Also recall the definition of $U_{i,k}$ and $V_{i,k}$ in the construction of G. Let x_i be a variable and let $k \in [d_i]$, then we have

Lemma 5. *Let σ be an assignment drawn according to its weight. Then for any $i \in [n]$ and any $k \in [d_i]$, except for probability* $\exp(-\Omega(m^{1/3}))$*, we have*

$$\big|N_\sigma(U_{i,k})\big| < \frac{1}{1+e} \cdot \Big(1 + m^{-1/3}\Big) \cdot \big|U_{i,k}\big|.$$

Proof. Pick any partial assignment σ' over vertices of G except those in $U_{i,k}$. Conditioned on σ', it is easy to see that the values of vertices in $U_{i,k}$ are independent. Each vertex in $U_{i,k}$ has $\Delta + \Delta'$ neighbors, each of which contributes a vertex weight of either β or 1 if it is assigned 0, and either 1 or γ if it is assigned 1. Since $\gamma < 1/\beta$, the total weight for assignment 1 is at least $\gamma^{\Delta+\Delta'} \geq e$ times the weight for assignment 0. The lemma follows from the Chernoff bound.

Given an assignment σ, we use σ_i to denote its restriction over vertices in H_i, and σ_{-i} to denote its partial assignment over vertices in G except H_i. We let $M_{\sigma_{-i}}(U_i)$ denote the subset of U_i whose unique neighbor outside of H_i is assigned 1. Using Lemma 5 and the union bound, we have

Corollary 3. *Let σ be an assignment drawn according to its weight. Except for probability* $\exp(-\Omega(m^{1/3}))$*, we have*

$$\big|M_{\sigma_{-i}}(U_i)\big| \geq \left(\frac{e}{1+e} - O\big(m^{-1/3}\big)\right) \cdot \big|U_i\big|, \quad \text{for all } i \in [n]. \tag{14}$$

It is also clear that Lemma 5 and Corollary 3 also hold for $V_{i,k}$ and V_i, respectively, by symmetry. Now we are ready to prove Lemma 2. Let $\sigma = (\sigma_i, \sigma_{-i})$ be an assignment drawn from this distribution. Recall the definition of Σ below (9). Then by Corollary 3 we have

$$\Pr\big[\sigma \notin \Sigma\big] \le \exp(-\Omega(m^{1/3})) + \Pr\big[\sigma \notin \Sigma \;\big|\; \sigma_{-i} \text{ satisfies (14) for both } U_i \text{ and } V_i \text{ and for all } i \in [n]\big] \tag{15}$$

To prove an upper bound for (15) we fix σ_{-i} to be any partial assignment over the vertices of G except those of H_i, which satisfies (14) for both U_i and V_i. Then it suffices to prove that the sum of $\omega(G, \sigma)$ over all $\sigma \in \Sigma$ that are consistent with σ_{-i}, denoted by Z_1, is exponentially larger than the sum of $\omega(G, \sigma)$ over all $\sigma \notin \Sigma$ that are consistent with σ_{-i}, denoted by Z_2.

Let $\omega(\sigma_{-i})$ denote the product of the edge weights in σ_{-i} over all edges that have no vertex in H_i. By the definition of $Z_{a,b}(H)$ in (12), we have

$$Z_2 \le \omega(\sigma_{-i}) \sum_{a,b \in T_{d_i m}\,:\, a,b \ge \lambda} Z_{a,b}(H_i) \le \omega(\sigma_{-i}) \cdot (d_i m)^2 \cdot \exp\big(1.22 \cdot d_i m\big) \tag{16}$$

where the second inequality follows from (13). To get a lower bound for Z_1, let

$$L = \big|M_{\sigma_{-i}}(U_i)\big| \quad \text{and} \quad R = \big|M_{\sigma_{-i}}(V_i)\big|.$$

Consider all the σ that are consistent with σ_{-i} and $U_i(\sigma) = 0$. This gives us

$$Z_1 \ge \omega(\sigma_{-i}) \cdot \gamma^{\Delta' L} \cdot (1 + \gamma^{\Delta + \Delta'})^R \cdot (\beta^{\Delta'} + \gamma^{\Delta})^{d_i m - R}.$$

Plugging in $\gamma^{\Delta+\Delta'} \ge e$, $\gamma^{\Delta} \ge e^{7999/8000}$ and the lower bound in (14), we get

$$Z_1 \ge \omega(\sigma_{-i}) \cdot \exp\big(1.22897 \cdot d_i m\big)$$

and the lemma follows from (16).

References

1. Lovász, L.: Operations with structures. Acta Mathematica Hungarica 18, 321–328 (1967)
2. Hell, P., Nešetřil, J.: Graphs and Homomorphisms. Oxford University Press (2004)
3. Dyer, M., Greenhill, C.: The complexity of counting graph homomorphisms. In: Proceedings of the 9th International Conference on Random Structures and Algorithms, pp. 260–289 (2000)
4. Bulatov, A., Grohe, M.: The complexity of partition functions. Theoretical Computer Science 348(2-3), 148–186 (2005)
5. Goldberg, L., Grohe, M., Jerrum, M., Thurley, M.: A complexity dichotomy for partition functions with mixed signs. SIAM Journal on Computing 39(7), 3336–3402 (2010)
6. Cai, J.Y., Chen, X., Lu, P.: Graph homomorphisms with complex values: A dichotomy theorem. In: Proceedings of the 37th Colloquium on Automata, Languages and Programming (2010); To appear in SIAM Journal on Computing

7. Dyer, M., Goldberg, L., Paterson, M.: On counting homomorphisms to directed acyclic graphs. Journal of the ACM 54(6) (2007)
8. Cai, J.Y., Chen, X.: A decidable dichotomy theorem on directed graph homomorphisms with non-negative weights. In: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, pp. 437–446 (2010)
9. Bulatov, A.: The complexity of the counting constraint satisfaction problem. In: Proceedings of the 35th International Colloquium on Automata, Languages and Programming, pp. 646–661 (2008)
10. Dyer, M., Richerby, D.: On the complexity of #CSP. In: Proceedings of the 42nd ACM Symposium on Theory of Computing, pp. 725–734 (2010)
11. Cai, J.Y., Chen, X., Lu, P.: Non-negatively weighted #CSP: An effective complexity dichotomy. In: Proceedings of the 26th Annual IEEE Conference on Computational Complexity, pp. 45–54 (2011)
12. Cai, J.Y., Chen, X.: Complexity of counting CSP with complex weights. In: Proceedings of the 44th ACM Symposium on Theory of Computing (2012)
13. Jerrum, M., Sinclair, A.: Polynomial-time approximation algorithms for the ising model. SIAM Journal on Computing 22(5), 1087–1116 (1993)
14. Goldberg, L., Jerrum, M., Paterson, M.: The computational complexity of two-state spin systems. Random Structures and Algorithms 23(2), 133–154 (2003)
15. Dyer, M., Frieze, A., Jerrum, M.: On counting independent sets in sparse graphs. SIAM Journal on Computing 31, 1527–1541 (2002)
16. Weitz, D.: Counting independent sets up to the tree threshold. In: Proceedings of the 38th Annual ACM Symposium on Theory of Computing, pp. 140–149 (2006)
17. Sinclair, A., Srivastava, P., Thurley, M.: Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. In: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (2012)
18. Li, L., Lu, P., Yin, Y.: Approximate counting via correlation decay in spin systems. In: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (2012)
19. Li, L., Lu, P., Yin, Y.: Correlation decay up to uniqueness in spin systems. arXiv:1111.7064 (2011)
20. Bandyopadhyay, A., Gamarnik, D.: Counting without sampling: Asymptotics of the log-partition function for certain statistical physics models. Random Structures and Algorithms 33, 452–479 (2008)
21. Sly, A.: Computational transition at the uniqueness threshold. In: Proceedings of the IEEE 51st Annual Symposium on Foundations of Computer Science (2010)
22. Galanis, A., Ge, Q., Štefankovič, D., Vigoda, E., Yang, L.: Improved inapproximability results for counting independent sets in the hard-core model. In: Proceedings of the 15th International Workshop on Randomization and Computation, pp. 567–578 (2011)
23. Cai, J.Y., Chen, X., Guo, H., Lu, P.: Inapproximability after uniqueness phase transition in two-spin systems. University of Wisconsin, Madison CS Technical Report (2011), `http://digital.library.wisc.edu/1793/61488`
24. Mossel, E., Weitz, D., Wormald, N.: On the hardness of sampling independent sets beyond the tree threshold. Probability Theory and Related Fields 143, 401–439 (2009)
25. Håstad, J.: Some optimal inapproximability results. Journal of the ACM 48, 798–859 (2001)